July 15, 1996 • Volume II, Issue 6

COMPUTERWORLD

· Leadership Series ·

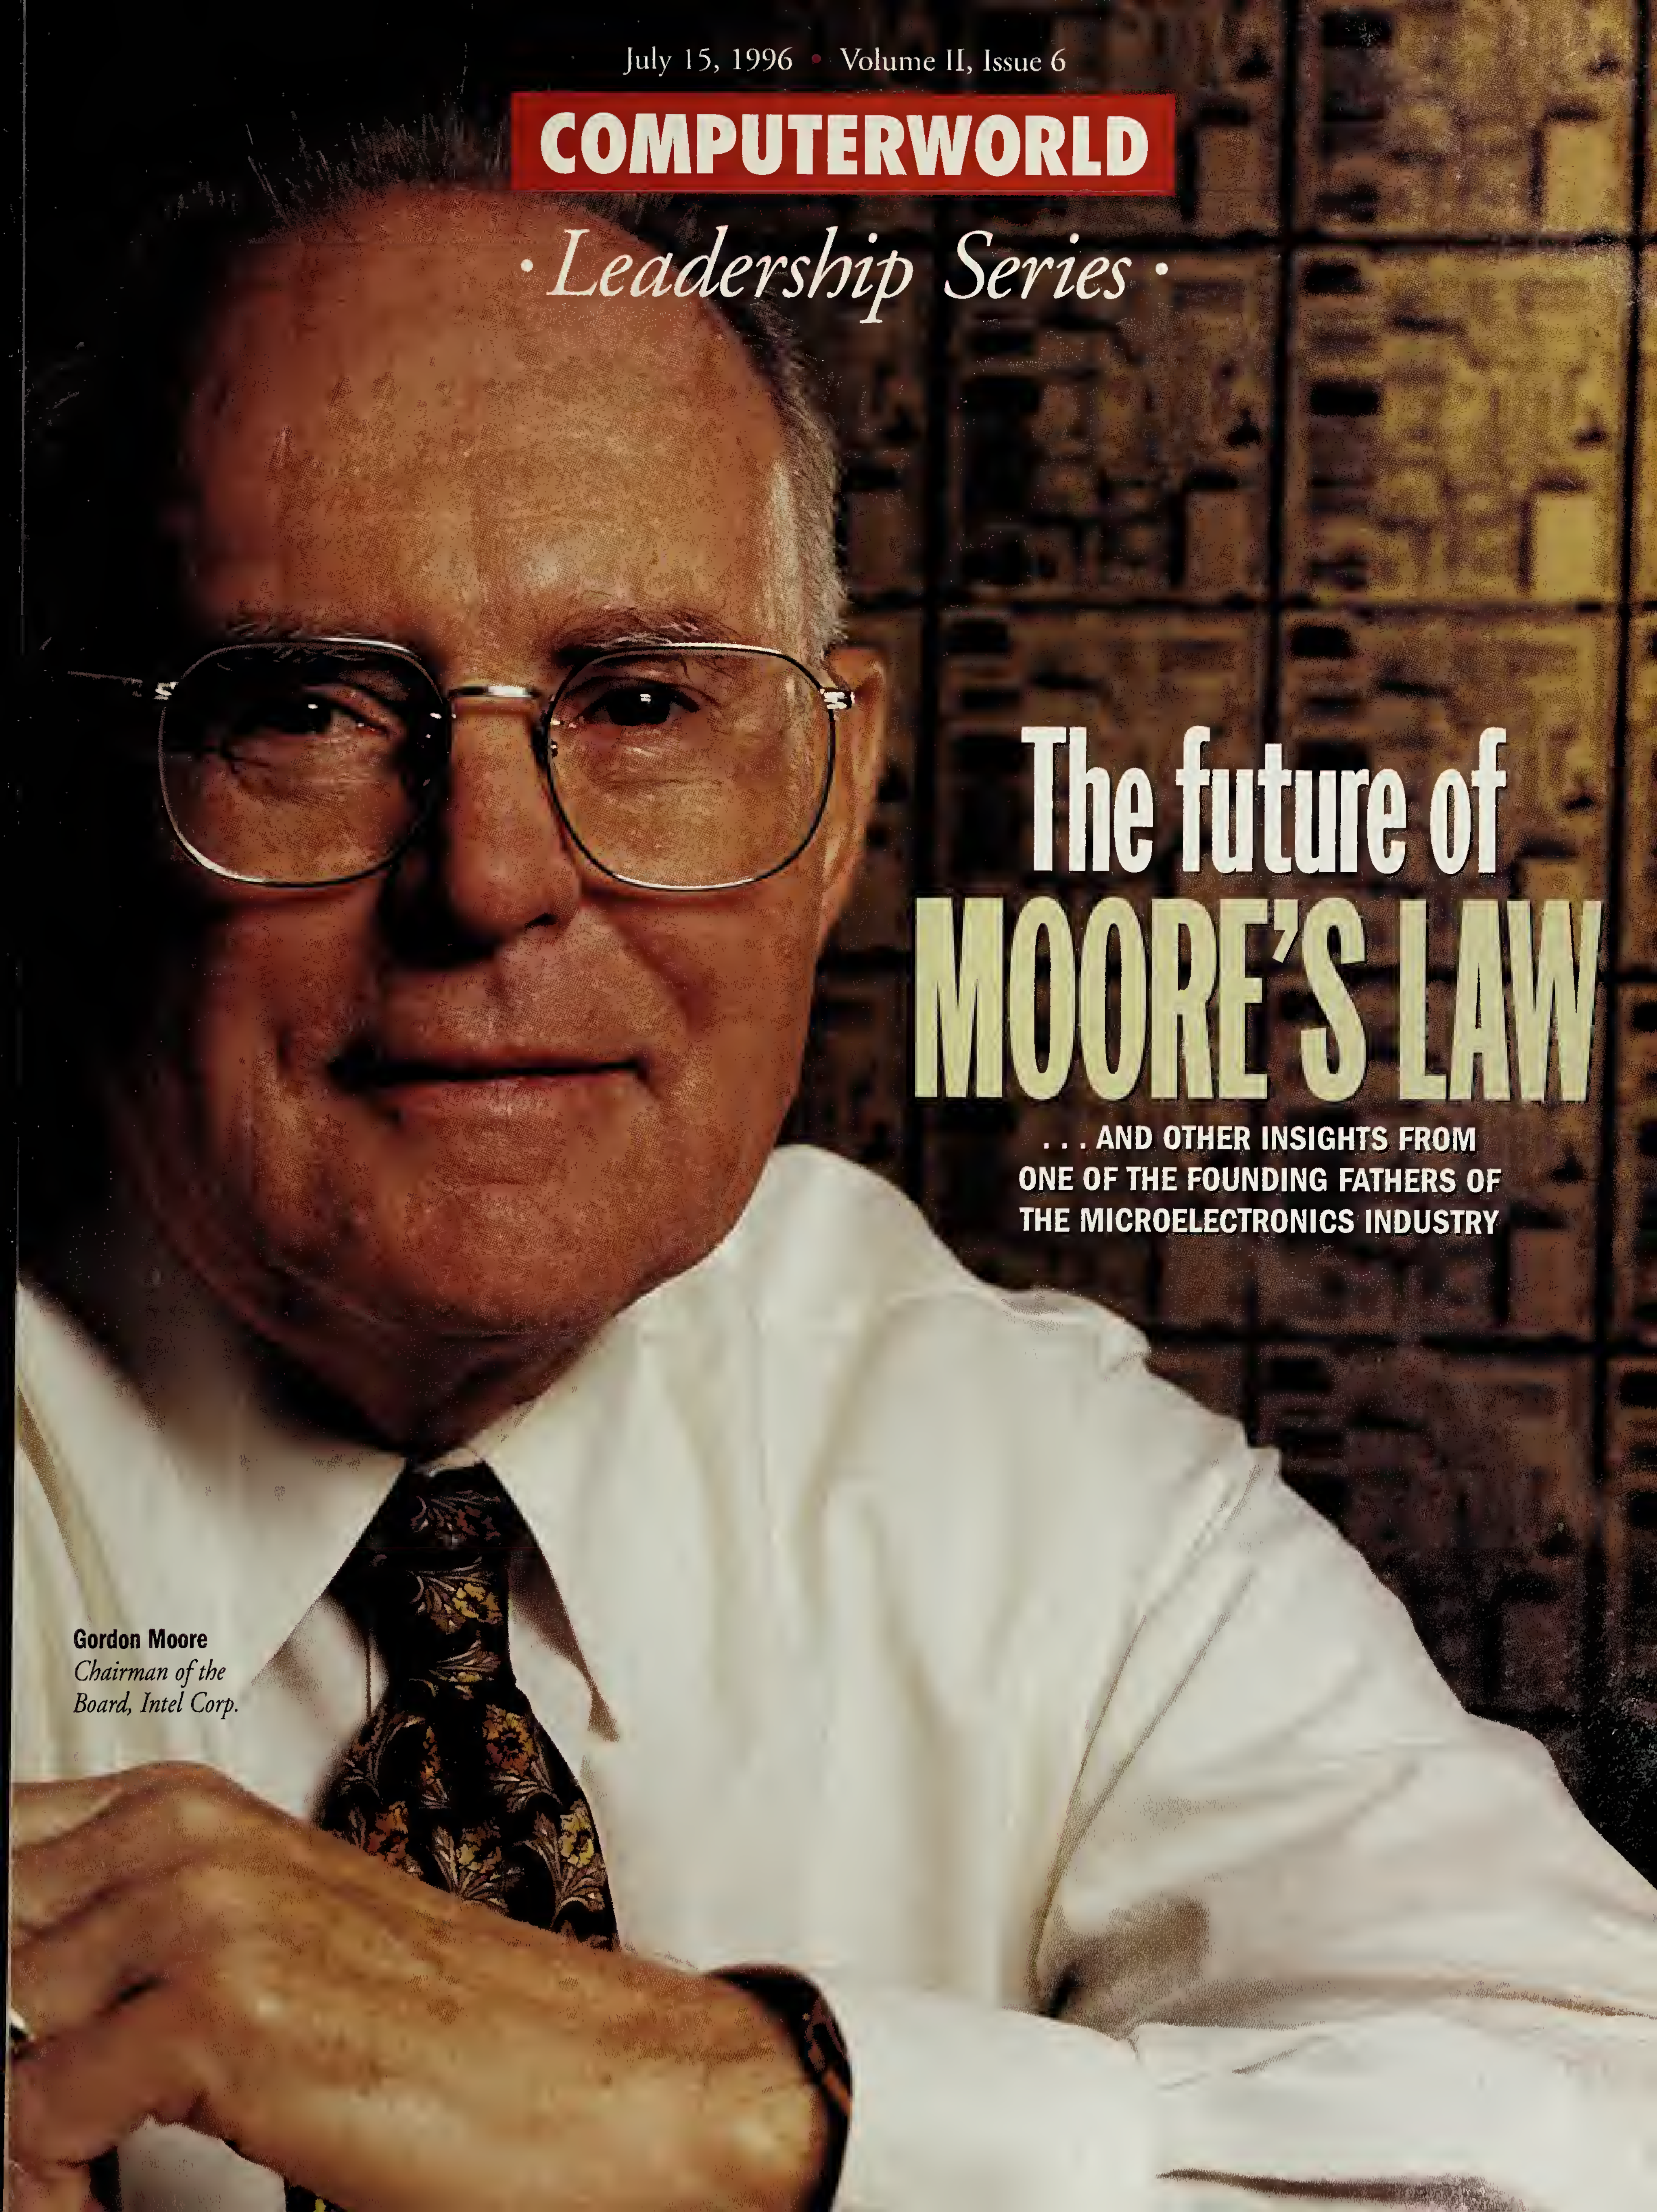

The future of MOORE'S LAW

. . . AND OTHER INSIGHTS FROM ONE OF THE FOUNDING FATHERS OF THE MICROELECTRONICS INDUSTRY

Gordon Moore
Chairman of the Board, Intel Corp.

Can Moore's Law continue indefinitely?

Probably not, says Intel Chairman and Co-founder Gordon Moore. Technical and business limitations will soon present real problems to the "law" that states chip density doubles every 18 to 24 months. Falling off the curve has implications for everyone in the silicon food chain, including IS managers who count on rapid systems upgrades to drive business.

At the Computerworld-Smithsonian Monticello Lectures on May 1, Moore spoke on the past, present and future of the microchip, the electronics industry and Moore's Law. This article is excerpted from his lecture and from an interview with Leadership Series Editor Bruce Rayner.

There is nothing natural about Moore's Law. It emerged from the work of people developing dramatic new technologies, products and processes. And there is nothing sacred about the 18-to-24-month time period between generations of microchips. The relentless pace of chip development has to do with many different factors. Among them: having access to a lot of capital, moving research and development close to the manufacturing process, creating open organizations where innovation can flourish, accepting change as a constant and, of course, having a healthy dose of luck. Indeed, these are the factors that have shaped Silicon Valley over the last four decades.

But there is another, more fundamental driver of Moore's Law, according to the man himself. "More than anything, once something like this gets established, it becomes more or less a self-fulfilling prophecy. The Semiconductor Industry Association puts out the technology road map, which continues this generation [turnover] every three years. Everyone in the industry recognizes that if they don't stay on essentially that curve they will fall behind. So it sort of drives itself. Until we get to the point where, for some reason or another, no matter how hard we try we can't do it anymore."

What will stop the silicon freight train? Moore, who in 1965 first identified the relationship between time and the number of transistors on a microchip, has two fundamental concerns about staying on the curve. First, there are business limitations. As chip densities rise, the cost of production rises almost exponentially. Second, there are real physical limitations. Problems arise as chip design gets down to the atomic level.

Will a new technology, or combination of technologies, provide a solution? Moore doesn't think so. "It is always presumptuous to say 'forever,' but I don't see anything that is going to come along that will replace [microelectronics] technology.

"Maybe I'm just getting old, but my view is that the technology that's developed around silicon was a general-purpose technology for making microstructures of materials. I think it's extremely unlikely that some kind of biological assembly process, to pick an example, is going to be able to make the same kind of products, and therefore, re-

Photo illustration by John Baxter

ally be a substitute technology."

If that's the case, how much longer can the semiconductor industry sustain exponential growth in the complexity of successive generations of integrated circuits?

"I think much of the rate of progress can be expected to continue for at least a few more generations. Three generations of technology at three years per generation is about a decade. So I can see us staying on roughly the same curve that long."

Business limitations

"What concerns me more than approaching physical limits — and there are real physical limits as eventually the atomic nature of matter really starts to bite — is the fact that the cost of the technology seems to be growing exponentially as well. This doubling of the cost about every generation is something that is going to be hard to continue. It will be hard to stay on the same kind of curve that we've been on previously."

Cost of equipment. "When Intel started [in 1968], we raised $3 million. That let us equip our first factory, develop new technologies and develop our first products and get them to market. Now one of our larger factories that does just the wafer processing — it doesn't even do the assembly and test — costs about $2.4 billion when they're full of equipment. We also have smaller factories that cost $1.5 billion. [These new factories] don't last very long because of the [rapid evolution of] technology.

"Back in 1968, a piece of manufacturing equipment — a bank of furnaces or a photo lithography tool or whatever — cost about $12,000. It varied a little bit, but that's a pretty good approximation. The number now is about $6 million — a 500-fold increase in the cost of a piece of equipment.

"Intel's revenue per employee has grown from about $64,000 [in the mid-1980s] to $437,000 last year, a sixfold improvement. But it's a darn good thing it did because our capital investment per employee grew from $68,000 to $274,000, or about fourfold. That's a pretty darn capital-intensive business.

"If you extrapolate the history of the last few decades, the cost of a factory has about doubled every technology generation, and it doesn't show any sign of rounding off. If you double that for even two more generations, you arrive at $10 billion factories. That is a rate of increase that is awfully hard to swallow.

"I don't think we can do that for two more generations. Of course, two generations ago, I wouldn't have considered the $2.5 billion factory viable, either. We are getting to the point where

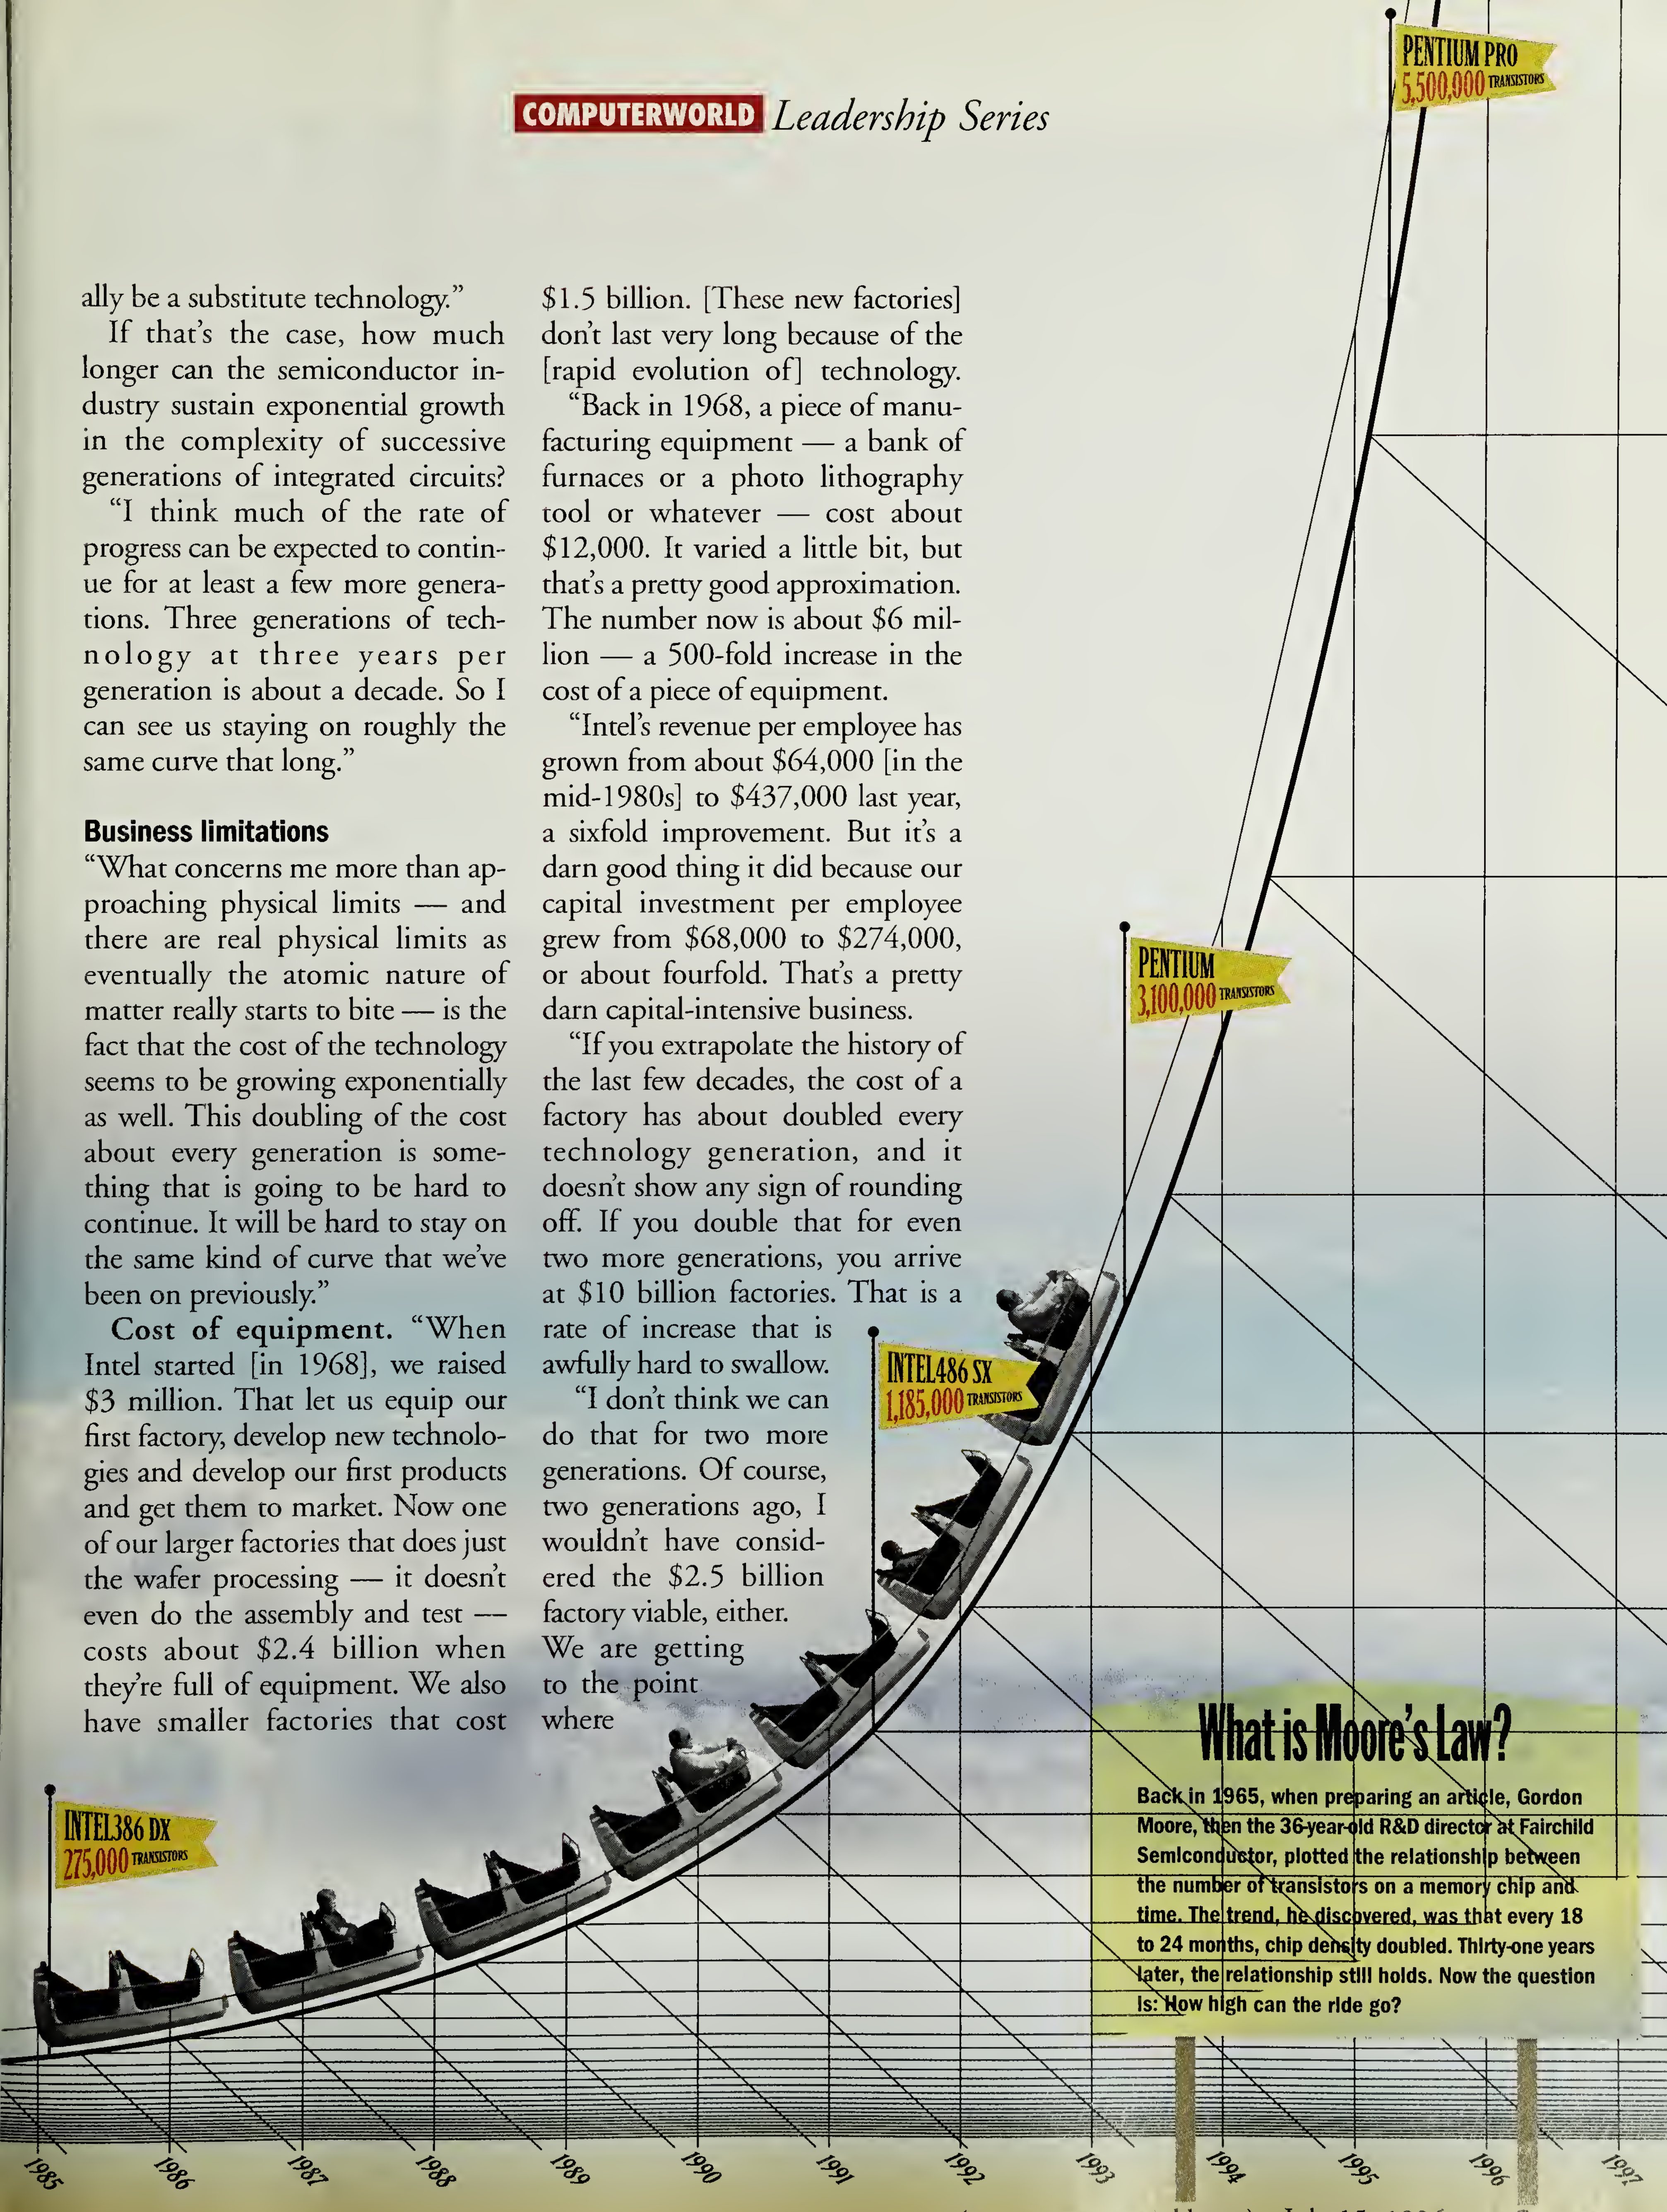

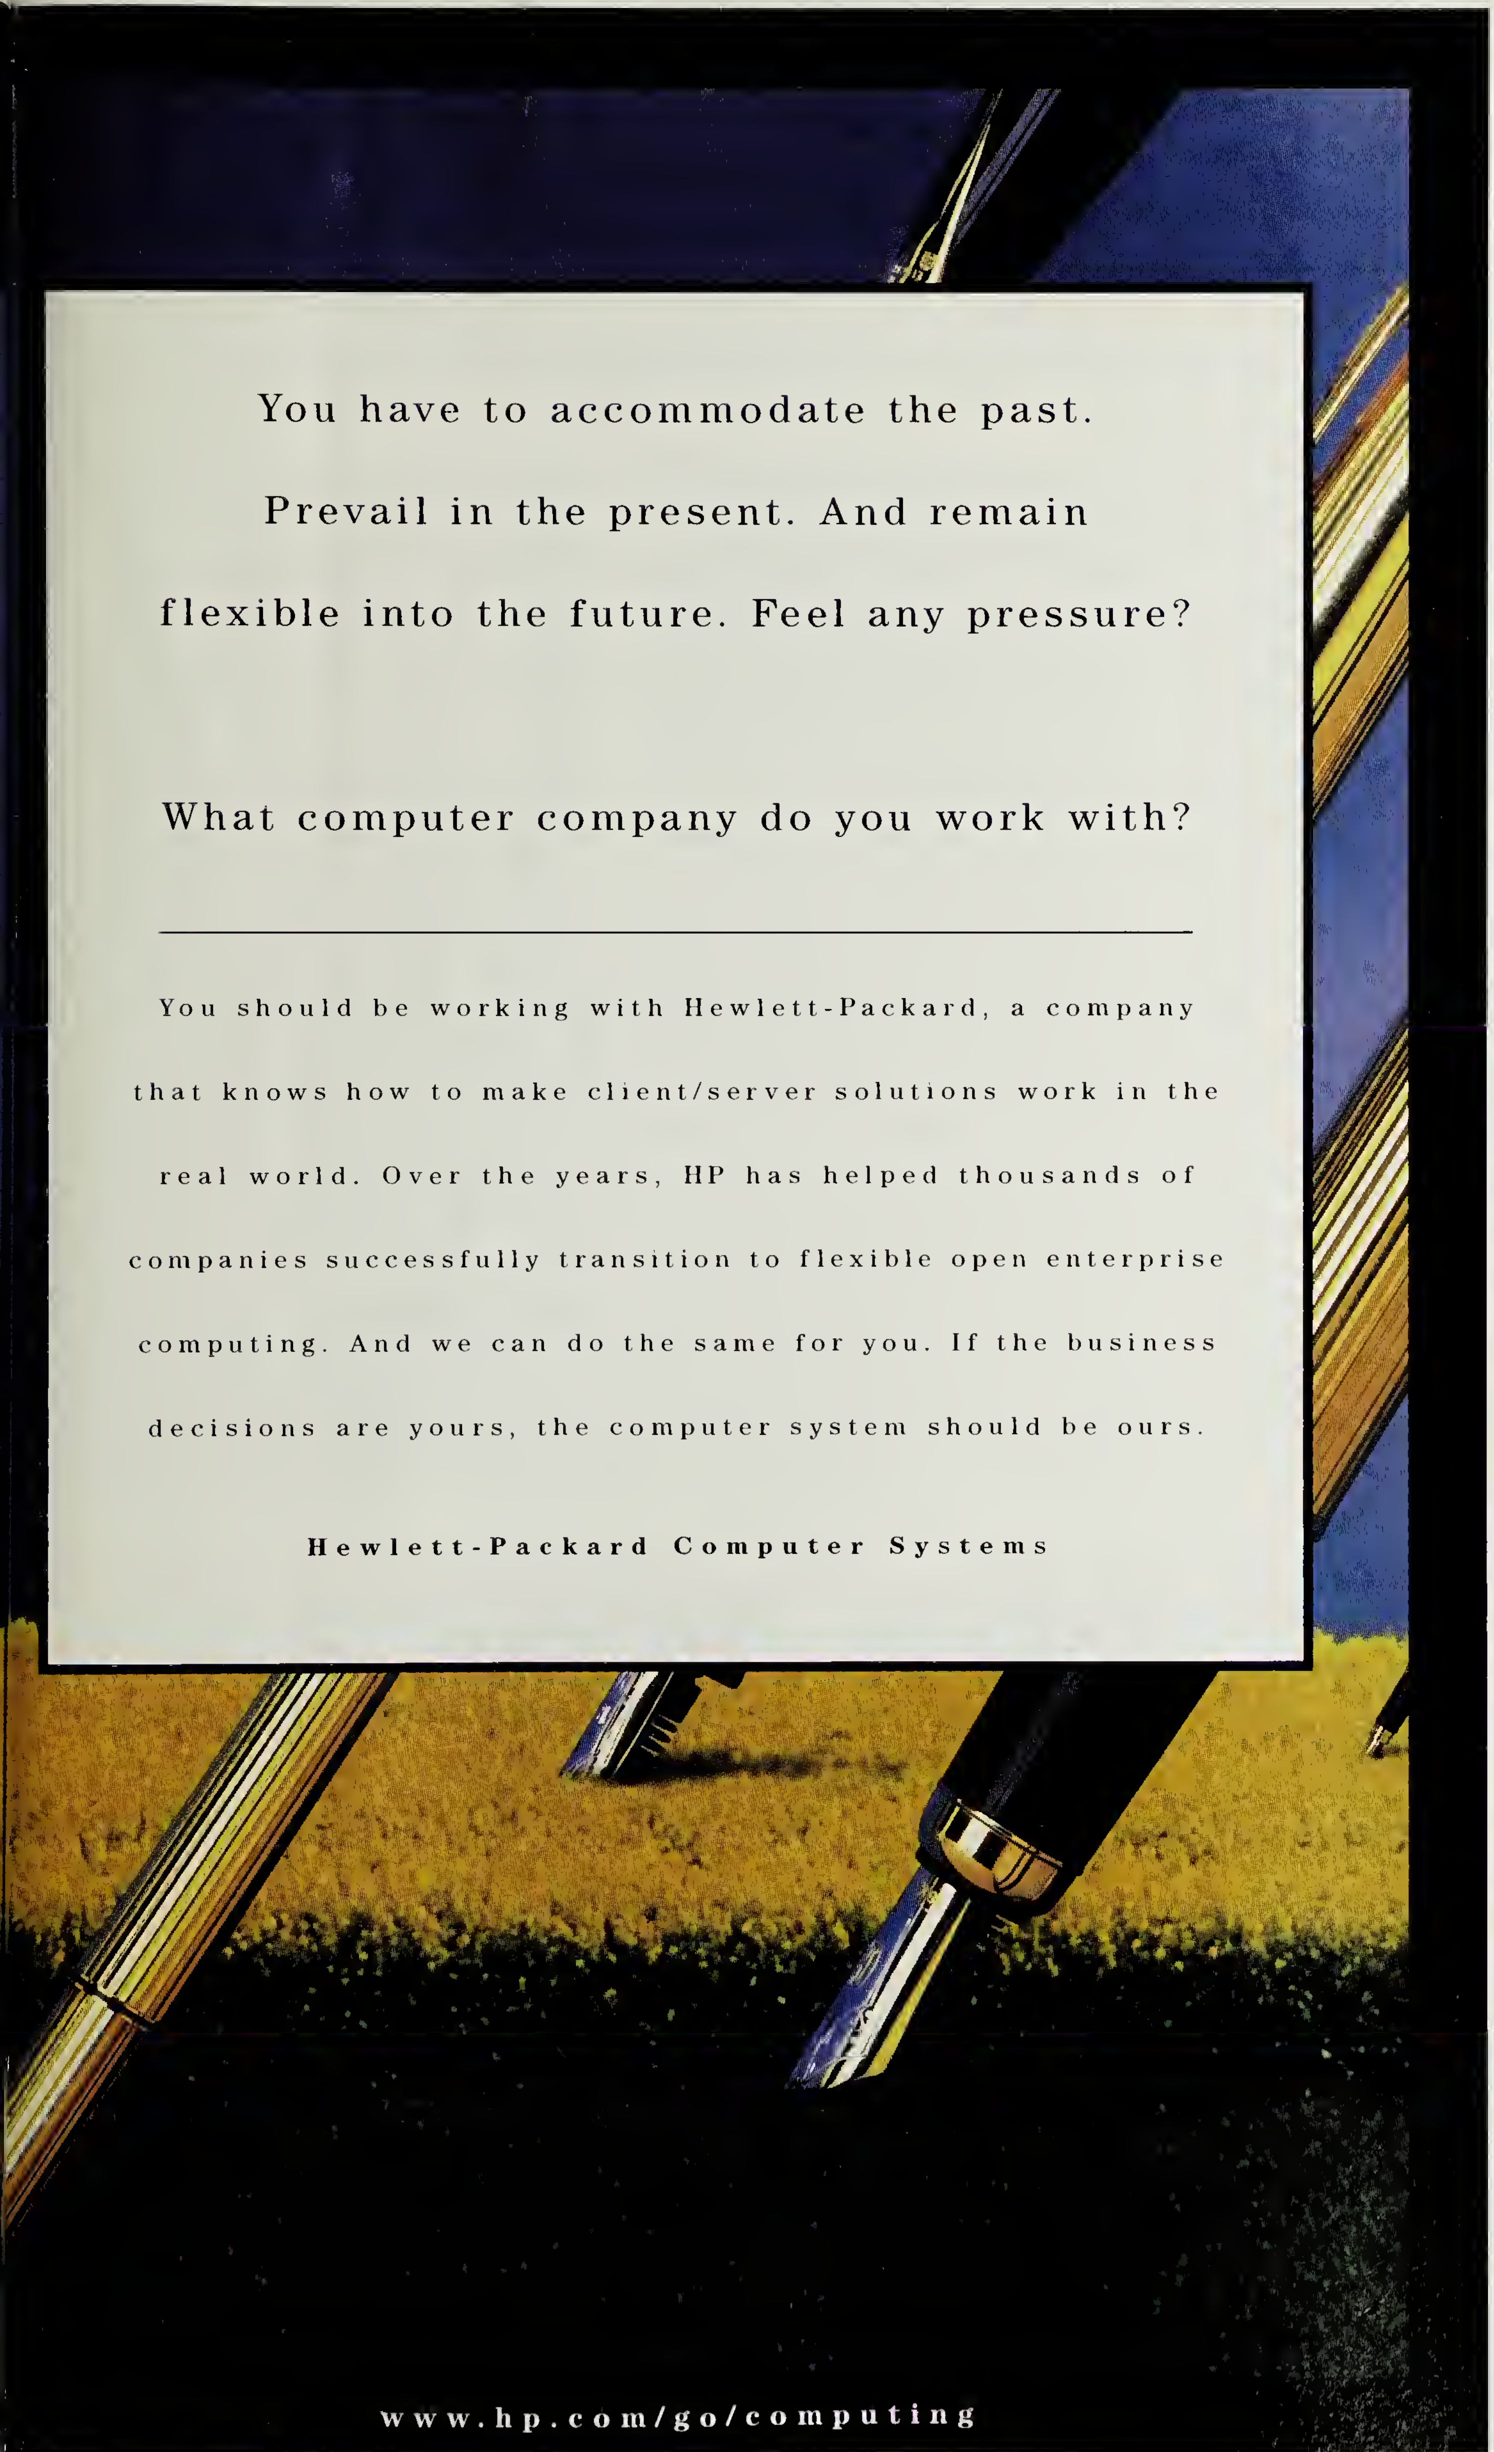

HEWLETT® PACKARD

After sending the first simple integrated circuit into production in 1961, Moore called his lab staff together to discuss what technology they'd tackle next. "We had no appreciation that we had just turned over the first stone in this huge quarry."

the cost of the factory, the depreciation cost and the cost of the investment is an increasingly large fraction of the total cost of building the integrated circuit. It is getting up to the percentage now that it is certainly not clear to me that we can keep doing it.

"A lot of the smaller companies, even some larger companies, are taking advantage of factories built by others. There are factories built in Asia whose charter from the beginning was to supply contract manufacturing.

"But I don't anticipate the major players moving in that direction. You might avoid the capital investment, but there are a lot of disadvantages to depending on other people's capacity. The only way you can be sure you have leading-edge capacity is if you build it yourself."

Cost of labor. "Intel's first microprocessor, [the 4004], which came out in 1971 and had about 2,200 transistors on it, had a gang of about 10 people designing it. Today, the generation we're working on, instead of 2,200 transistors, has about 10 million transistors, or about a 5,000-fold increase in the complexity of the chip if you measure it by the number of transistors. The size of the design team has gone from 10 people to 400 to 500 people, and the costs have gone up at least proportionally. If you plot it, it is pretty near exponential. Unfortunately, I see [this trend] continuing."

Physical limitations

The other major challenge the semiconductor industry faces is the physical dimensions of the chips themselves. As the number of tran-

The $15 million dollar watch

It may be a microprocessor powerhouse, but Intel missed the mark on some key opportunities along the way, admits Moore:

DIGITAL WATCHES. "Our initial objective was to identify areas where we could make complex silicon chips in large volume: Memories being the first one, microprocessors being the second one. We thought digital watches was the third one.

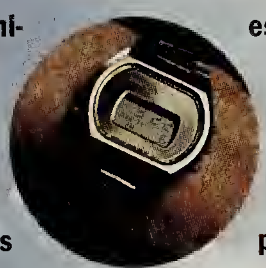

"Intel was the first company with a liquid crystal digital watch on the market [in the early 1970s]," says Moore, "and we were also one of the first ones out of the business."

The total investment in Intel's Microma venture? About $15 million, Moore estimates. He still straps on his vintage digital watch every day, keeping it as a reminder that not all business ventures will be as lucrative as microprocessors. "It has an alarm in it that wakes me up every morning."

THE HOME PC. The personal computer "was an opportunity that was presented to me [in the early 1970s], well before Apple [developed its first

Photography by Andy Freeberg

sistors rises, so does the complexity of the design, the size of the transistors, manufacturability and so on.

Transistor dimensions. "It looks like the transistors themselves face some problems with minimum dimensions on the order of 0.05 micron." Right now, the industry is moving toward 0.25 micron technology for mass production.

"If you can operate at increasingly lower voltage levels you'd be OK, but people are concerned with going below something like one volt. I think that practically now, the people in the industry would generally agree that a minimum dimension of around 0.05 microns starts to cause some fundamental problems with devices.

"The [integrated circuit] industry tends to go in steps of about a factor of 0.7. So if we are looking at 0.25 micron as the next generation, then a few years later we'll be at 0.18 micron [approximately 0.25 micron times 0.7]; then a few years after that, 0.13 micron [approximately 0.18 times 0.7]. If you stayed on that path, we've got another two or three generations after 0.13 before we get to 0.05 micron." If each generation takes about three years, then it will be 12 years before the industry faces the 0.05 micron barrier.

X-ray lithography. "Now, there are some intermediate barriers that we have to contend with. One of them is ordinary optical lithography." (Optical lithography is the technology that chip makers use to produce the device pattern on the silicon wafer.)

"Ordinary optical lithography systems can go down to 0.13 micron. But to get below that, you move into X-ray region. X-ray or 'extreme ultraviolet' lithography is a major change for the industry. It will require completely new equipment. So X-ray lithography has got a lot of investment ahead.

"I think to avoid [technical] problems, people will come up with a hybrid approach to lithography, where you only use [X-ray lithography] for the finer structures in a small portion of the device and then [use] regular photo lithography for the rest [of the chip]. I think there's going to be some kind of a discontinuity, though."

What next?

Barring a replacement technology, Moore's Law has until about 2010 before it comes up against a brick wall, Moore says. What then? Will product development begin to slow? Will the cost of successive generations of technology rise instead of fall, as it has over the past three decades?

A lot depends on technology innovations over the next few years. But it wouldn't hurt for business leaders to start thinking about the consequences of the end of Moore's Law. ◆

Gordon E. Moore, *Ph.D., 67, is chairman of the board at Intel Corp. A pioneer in the microelectronics industry and co-founder of Intel in 1968, he received the National Medal of Technology in 1990 from President George Bush.*

---

model]. An Intel engineer came up to me and suggested that we could make a computer for the home using the microprocessors that we were doing then. But the only application he could suggest was the housewife with her recipes.

"I couldn't see my wife sitting at the stove poking numbers into a PC, and so we decided that it's not a product that would ever go anyplace."

DRAMs. "We missed some opportunities by poor activities: The dynamic RAM business, for example, which was the product that really got [Intel] going initially. [In the early 1980s], we kind of stumbled on a couple of subsequent [DRAM] generations and lost our leadership position.

"And eventually in the [electronics] industry depression in 1985, where we were positioned again to jump back into a leadership position, we looked at the investment that was required and decided we would have to build a couple of new plants. At that time, [the investment] would have been about $400 million. And since the whole industry was losing money in that business, we didn't see it as a good return.

"So in spite of having done the development of both the technology and the product, we abandoned DRAMs."

Computerworld Leadership Series: Editor: Bruce Rayner; Managing Editor: Joyce Chutchian-Ferranti, Anne McCrory; Designer: Mary Beth Welch. For article submissions, contact Paul Gillin at (508) 620-7724, or E-mail him at paul_gillin@cw.com. For previous Leadership Series articles, call (508) 820-8125.


THANKS.
Let's be honest. Having our NetServers showered with accolades is a kick.
Besides, it shows how serious we are about building quality network systems.
For detailed specs, visit http://www.hp.com/go/netserver. Or call 1-800-533-1333, Ext. 1690.
HP NETSERVERS
NetServer 5/100
LAN
PRODUCTS OF THE YEAR
PC MAGAZINE EDITORS' CHOICE
March 26, 1996
HP NetServer 5/133 LH
PC Computing
BEST
Computer Reseller News
CHANNEL CHAMPION
1995
EDITOR REFUSES TO GIVE IT BACK
NETWORK COMPUTING
MVP
PC Computing
intel inside
pentium PROCESSOR